

Contact Plan for Time-Variant Routing

draft-blanchet-tvr-contactplan

Marc Blanchet
marc.blanchet@viagenie.ca
IETF 116 Yokohama

Rationale

- Space communications are scheduled based on time windows to communicate. Therefore, a contact plan is needed for communications stacks to know when to expect a contact.
- Motivation:
 - currently developing a Bundle Protocol stack in Swift for Apple platforms (and Linux): Aim is to get BP apps on an (Astronaut) iPhone. (Android is next).
 - Questions came up: data model/file format for contact plans; policies regarding holding/dropping/forwarding bundles
- Looked at available BP stack implementations for contact plans: everyone has a different format for specifying contact plan.: for example: ION: transaction based (add/delete/edit); ud3tn: line based with various separators; hdtv: json.
- Purpose of the draft:
 - specify a standardized data model and file format for contact plan.
 - While coming from BP, generalized it for IP.
 - Enables contact plan engineers to provide a single format, whatever the implementation being used.
 - Enables comparison, diffs, revisions, archiving; tools to display, manage, design, ...

An Example

```
{
  "type": "tvrContactPlan",
  "version": 1,
  "lastUpdated": "2025-01-17T23:20:50Z",
  "contacts": [
    {
      "id": "f81d4fae-7dec-11d0-a765-00a0c91e6bf6",
      "family": "ip4",
      "source": "192.0.2.0",
      "destinations": ["198.51.100.0/24", "203.0.114.0/28", "192.0.3.1"],
      "nextHop": "203.0.113.1",
      "startTime": "1985-04-12T23:20:50Z",
      "stopTime": "1985-04-13T14:12:48Z",
      "bandwidth": 1000000,
      "latency": 30000
    },
    {
      "id": "f81d4fae-abcd-efgh-a765-00a0c91e6b88",
      "family": "ip6",
      "source": "2001:db8::1",
      "destinations": ["2001:db8:abcd::/48"],
      "nextHop": "2001:db8:3::1",
      "startTime": "2030-04-12T23:20:50Z",
      "stopTime": "2031-04-13T14:12:48Z",
      "bandwidth": 10000000,
      "latency": 300000
    }
  ],
}
```

An Example

```
{
  "id": "659e4fae-7dec-11d0-a765-00a0c91e6b04",
  "family": "dtn",
  "source": "dtn://ud3tn2.dtn/",
  "destinations": ["dtn://18471/", "dtn://81491/"],
  "nextHop": "203.0.113.1",
  "startTime": "1985-04-12T23:20:50Z",
  "stopTime": "1985-04-13T14:12:48Z",
  "bandwidth": 1000000,
  "latency": 30000
},
{
  "id": "f81dab43-7dec-e8a2-a765-00a0c91e6bf6",
  "family": "ipn",
  "destinations": "ipn:5.34",
  "nextHop": "ipn:7.43",
  "startTime": "1985-04-12T23:20:50Z",
  "stopTime": "1985-04-13T14:12:48Z",
  "bandwidth": 1000000,
  "latency": 30000
}
]
```

Comments Received/TODO

- Use Yang (Dean Bogdanovic, Tony Li): will do in next rev. Dean has offered to help
- Unnumbered links (Tony Li): good point. Will look into it and do something in next rev.
- Why IP (Tony Li): some use cases where delays is « not too much » but still requires time-variant routing.
- Top-level id field and index for each « record » (Tony Li):
 - pros and cons on each approach.
 - Having an id on each record enables editing, removing, adding new, ... without changes to all indexes.
 - One implementation actually does not have a file format per se(internal), but instead transactions: add/delete/edit so id on each record is more compliant with this implementation.
 - But top-level id field is also good for identifying the whole contact plan.
 - Maybe both levels have ids. Not yet sure what is best.

Comments Received/TODO

- Support empty list of contacts(Tony Li): ok. Will do in next rev.
- Why a list of destinations(Tony Li): BP does not really have a way to specify a prefix as CIDR in IP. Therefore, it has to be a list of node ids. Moreover, if various destinations have the same characteristics, then it makes sense to use the same record and just specify a list of destinations.
- Why start and stop time mandatory? (Tony Li): in BP implementations it is, but I agree it may be unspecified. Proposing to make mandatory at least one? Or do we want to allow to not specify any time? What would be the semantic then? (Always?)
- Bandwidth and latency, why? (Tony Li): in BP implementations, bandwidth is specified. Latency is to help transport for hints, but I know it is a layer violation: just wanted to avoid « another » contact plan for transport « just » for latency specifications.
- Convergence layer syntax (Felix Walter): yes need to do it. Next rev.
- Needed (Felix Walter)

Comments Received/TODO

- No need for this file format, implementation decisions (Scott Burleigh). Disagree, exchange, archiving, distribution, tools. At least, the data model is useful. Decision to use this file format is up to the implementor.
- Specify more properties such as probability of contact (Felix Walter): used in some BP implementations. Used? Maybe define a registry of properties so that it could be extended.
- Wildcard for DTN URI: used in ION.

Next Steps

- Thanks to early reviewers: Scott Burleigh, Felix Walter, Tony Li, Dean Bogdanovic
- WG Adoption
- Specification: [draft-blanchet-tvr-contactplan](#)
- Looking for more comments
- marc.blanchet@viagenie.ca